arXiv:1607.06849v1 [stat.ME] 22 Jul 2016

Reciprocal Graphical Models for Integrative Gene Regulatory Network Analysis

Yang Ni[1], Yuan Ji[2], and Peter Müller[3]

[1]Department of Statistics and Data Sciences, The University of Texas at Austin

[2]Program for Computational Genomics & Medicine, NorthShore University HealthSystem, Department of Public Health Sciences, The University of Chicago

[3]Department of Mathematics, The University of Texas at Austin

Abstract

Constructing gene regulatory networks is a fundamental task in systems biology. We introduce a Gaussian reciprocal graphical model for inference about gene regulatory relationships by integrating mRNA gene expression and DNA level information including copy number and methylation. Data integration allows for inference on the directionality of certain regulatory relationships, which would be otherwise indistinguishable due to Markov equivalence. Efficient inference is developed based on simultaneous equation models. Bayesian model selection techniques are adopted to estimate the graph structure. We illustrate our approach by simulations and two applications in ZODIAC pairwise gene interaction analysis and colon adenocarcinoma pathway analysis.

Keywords: Simultaneous equation models; Markov equivalence; directed cycles; feedback loop; multimodal genomic data.

1 Introduction

In this paper, we develop a reciprocal graphical model (RGM) to infer gene regulatory relationships and gene networks. This includes in particular directed edges without time course or interventional data. RGMs allow for undirected edges, directed edges and directed cycles and therefore are ideally suited for modeling regulatory relationships including feedback loops. Exploiting genomic data from multiple modalities/platforms, we are able to determine the directionality of certain regulatory relationships, which would be otherwise indistinguishable due to Markov equivalence. Such inference about directionality becomes possible because basic biology fixes the directionality for some edges, for example, between DNA methylation and gene expression of the same gene. Conditioning on such known directionality enables us to infer directionality for other edges. Statistically, the class of probability models determined by RGMs is strictly larger than the class of probability models determined by directed acyclic graphs (DAGs) and Markov random fields (MRFs). Computationally, the connection of RGMs with simultaneous equation models (SEMs) facilitates computation-efficient implementation of full posterior inference.

Most recent graphical model approaches in biostatistics and bioinformatics are restricted to DAGs (Stingo et al., 2010; Yajima et al., 2015; Ni et al., 2015) and MRFs (Wang and West, 2009; Dobra et al., 2012; Green and Thomas, 2013; Mitra et al., 2013; Wang et al., 2013). These approaches use the conditional independence structure represented by the graphical models. The common use of MRFs and DAGs is due to mathematical tractability and easy computation, despite some inherent limitations. MRFs model undirected relationships between genes and do not account for the directionality of edges. However, biological interactions between genes are often asymmetric. For example, it is only possible for a regulator to regulate its targets, not vice versa. DAGs allow for directed edges but only as arbitrary factorization of a joint probability model. Also, DAGs explicitly prohibit directed cycles. However, feedback loops are quite common motifs and have key functional roles in many cellular processes such as regulating gene expressions and acting as bistable switches (Shin et al., 2010).

RGMs were first proposed by Koster (1996) but remain curiously under-used in the biostatistics and bioinformatics literature, with few exceptions. Zhang et al. (2005) developed

a hierarchical RGM to discover the relationship between cholesterol levels and pulmonary function in a longitudinal study. The graph structure is fixed and a generalized EM algorithm was developed to estimate the parameters. Telesca *et al.* (2012ab) use RGMs for latent variables that represent active/inactive proteins and differential vs. non-differential gene expression, respectively. However, the use of the RGM is restricted to representation and for convenient summaries. The actual inference model is based on the implied conditional independence structure only (after moralization). Also, they use RGMs with only directed edges, excluding possible undirected edges. By contrast, our use of RGMs is directly on the observed gene expressions and we use the RGM to build the probability model in a way that allows us to infer direction for directed edges.

Because different RGMs may determine equivalent conditional independence structure, inference algorithms that use only the implied conditional independence structure can not possibly distinguish such equivalent graphs based on observational data. In particular, the directionality of some regulatory relationships cannot be determined (examples are given in Section 2.2). In this paper, we propose a possible solution. We integrate different sources of genomic information including DNA copy number and DNA methylation in a way that allows us to determine the direction of some edges from first biological principles. This allows us then to identify the gene regulatory networks (GRN) that best fit the data, including inference on edge directions in some cases. That is, the additional genetic and epigenetic information together with fundamental biological knowledge can inform the directionality of the regulatory relationships between genes.

The proposed approach is motivated by two interesting applications. In the first application, we extend ZODIAC (Zhu et al., 2015) to incorporating directed edges between genes. ZODIAC is a free online tool (`http://www.compgenome.org/ZODIAC`) that provides (undirected) molecular interactions for about 200 million pairs of genes across different cancers based on multimodal genomic data from The Cancer Genome Atlas (TCGA). It can be used to explore gene regulations in cancer. With our contribution, ZODIAC will enable researchers to not only find out whether certain genetic interactions are likely to be present in cancer but also identify the sources (e.g. transcription factors) and targets (e.g. targeted genes) of the regulations. Our second application studies GRNs of colon adeno-

carcinoma (COAD). We focus on genes from the RAS-MAPK pathway which is critical in the initiation and progression of COAD (TCGA, 2012). Integrating copy number and methylation information, we find biologically meaningful gene interactions as well as novel regulatory relationships that need to be validated by further biological experiment.

Finally, for clarification and to avoid misunderstanding we note that our use of graphical models is quite different from inference for random graph models (Hoff, 2009; Goldenberg et al., 2010) when the graph itself is the data.

The rest of the article is organized as follows. We describe our model in Section 2. We present simulation studies in Section 3 and two applications in Section 4. Section 5 provides our closing discussion.

2 Model

2.1 Notation

A graph $\mathcal{G} = (V, E)$ consists of a set of *vertices* $V = \{1, \ldots, p\}$ and a set of *edges* E connecting these vertices. We consider both *directed* and *undirected* edges $E = E^d \cup E^u$ with $E^d \subseteq \{(i,j) \mid i,j \in V\}$ and $E^u \subseteq \{\{i,j\} \mid i,j \in V\}$ where the ordered pair (i,j) denotes a directed edge from vertex i to vertex j and $\{i,j\}$ denotes an undirected edge. We use the vertices to index a set of random variables, $\boldsymbol{Y} = \boldsymbol{Y}_V = (Y_1, \ldots, Y_p)^T$. For example, in our application Y_j represents gene expressions for gene j.

A *path* of length K is an ordered sequence $(i_0, \ldots, i_K)$ of distinct vertices except possibly $i_0 = i_K$ such that $\{i_{k-1}, i_k\} \in E$ or $(i_{k-1}, i_k) \in E$ for $k = 1, \ldots, K$. A path is called undirected if $\{i_{k-1}, i_k\} \in E$ for all $k = 1, \ldots, K$. A *path component* is a set of vertices that are all connected by an undirected path. A *reciprocal graph* (RG) is a graph $\mathcal{G} = (V, E)$ such that there are no directed edges between vertices in the same path component (Koster, 1996). Some examples of RG with four vertices are given in Figure 1a. The *boundary* of a vertex i is $\mathrm{bd}(i) = \{j \mid \{j,i\} \in E \text{ or } (j,i) \in E\}$ and the *boundary* of a subset $V_0 \subseteq V$ is $\mathrm{bd}(V_0) = \bigcup_{i \in V_0} \mathrm{bd}(i) \backslash V_0$. An anterior set is a subset $V_0 \subseteq V$ such that $\mathrm{bd}(V_0) = \emptyset$ and the smallest anterior set containing V_0 is denoted by $\mathrm{an}(V_0)$.

The Markov property (i.e. conditional independence relationships of $\boldsymbol{Y}$) of an RG

relies on the notion of *moralization*. To moralize a graph $\mathcal{G}$, we connect all vertices in the boundary of each path component of $\mathcal{G}$ by undirected edges and then replace all directed edges by undirected edges. The resulting moral graph is an undirected graph and is denoted by $\mathcal{G}^m$. Later we will use graph *separation* to introduce a global Markov property. In an undirected graph, two sets V_1 and V_2 are said to be separated by a third set V_3 if every path between V_1 and V_2 intersects V_3.

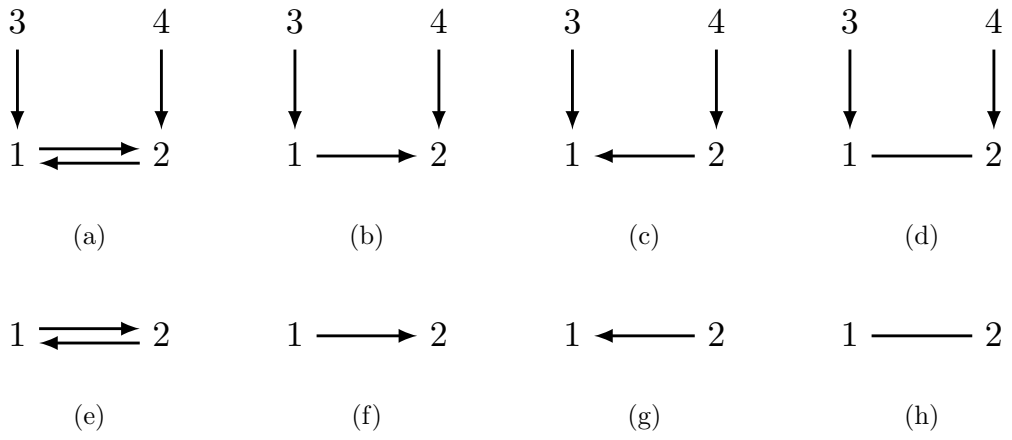

Figure 1: Graphs for illustration. Graphs in (a) – (d) are reciprocal graphs and imply different conditional independence. Graphs in (e) – (h) are Markov equivalent as they have imply the same Markov properties. In our applications, we will use nodes 1 and 2 to represent gene expression of two genes and nodes 3 and 4 to represent copy number (or methylation) of the same genes.

2.2 Reciprocal graphical models, Markov equivalence and integrative genomics

A graphical model is a mapping between a family of distribution and an underlying graph. In this paper, we focus on the class of RGs which by definition strictly contains MRFs and DAGs as subclasses.

The probability distribution of $\boldsymbol{Y}$ is said to be (global) *Markov* with respect to an RG $\mathcal{G}$ if $\boldsymbol{Y}_{V_1} \perp\!\!\!\perp \boldsymbol{Y}_{V_2} \mid \boldsymbol{Y}_{V_3}$ whenever V_3 separates V_1 and V_2 in $\mathcal{G}^m_{an(V_1 \cup V_2 \cup V_3)}$. RGMs can represent global Markov properties beyond the conditional independence structure that is encoded in MRFs and DAGs. RGMs are a strictly larger class of probability model than MRFs

and DAGs. For example, in Figure 1a, the only two conditional independence relationships are $3 \perp\!\!\!\perp 4$ and $3 \perp\!\!\!\perp 4 \mid 1, 2$. There is no MRF or DAG that encodes the same conditional independence relationships. More importantly for our application, RGMs are particularly useful for the construction of genomic networks because of the ability in modeling feedback loops which are not allowed by MRFs or DAGs. For example, in Figure 1a, gene 1 may regulate the expression of gene 2 while the status of gene 2 may reciprocally affect gene 1 through a feedback regulatory mechanism.

Two graphs are said to be *Markov equivalent* if they have the same Markov properties. For example, graphs 1e, 1f, 1g and 1h are Markov equivalent as they have the same Markov property, namely, no conditional independence assertion. Markov equivalence is indeed an equivalence relation and hence induces Markov equivalence class. Graphs within the same Markov equivalence class are usually nonidentifiable from observational data. In particular, observational data do not allow for inference on the direction of the edges in Figures 1e-1h.

However, with prior knowledge, sometimes we are able to distinguish the relationships between variables. In this paper, we introduce such prior knowledge by deliberately considering edges with (biologically) known direction if included. This allows for inference on other edges. We develop this approach based on integrating genomic information across different platforms and exploiting the central dogma of molecular biology that mRNA is produced by transcription from segments of DNA on which the copy number and methylation are measured, but the reverse processes are rare and biologically uninterpretable. For illustration, let vertices 1 and 2 in Figure 1 represent mRNA gene expressions and vertices 3 and 4 represent copy numbers of the corresponding genes and assume there exists dependence between copy number and gene expression for each gene and hence by the central dogma, there are directed edges from 3 to 1 ($3 \to 1$) and from 4 to 2 ($4 \to 2$). Now we are able to fully identify the relationship between mRNAs 1 and 2 because each of graphs 1a, 1b, 1c and 1d defines a distinct set of conditional independence relationships.

2.3 Simultaneous equation models

For a formal description of our approach, we still need the mapping from the RGM to a family of probability models for the observational data. Let $\boldsymbol{Y} = (Y_1, \ldots, Y_p)^T$ denote the

mRNA gene expressions for genes $1, \ldots, p$. Let $\boldsymbol{X} = (X_1, \ldots, X_{2p})^T$ be the set of DNA level measurements for genes $1, \ldots, p$ with X_{2i-1} and X_{2i} being the copy number and the methylation for gene i, respectively. We first state the simultaneous equation model (SEM) and will then introduce the mapping. An SEM for $\boldsymbol{Y}$ and $\boldsymbol{X}$ is given by

$$\boldsymbol{Y} = \boldsymbol{A}\boldsymbol{Y} + \boldsymbol{B}\boldsymbol{X} + \boldsymbol{E} \tag{1}$$

where $\boldsymbol{A} = (a_{ij}) \in \mathbb{R}^{p \times p}$ with zeros on the diagonal, $\boldsymbol{B} = (b_{ij}) \in \mathbb{R}^{p \times 2p}$, $\boldsymbol{E} = (\epsilon_1, \ldots, \epsilon_p)^T \sim \mathrm{N}_p(0, \boldsymbol{\Sigma})$, $\boldsymbol{X} \sim \mathrm{N}_{2p}(\boldsymbol{0}, \boldsymbol{\Psi})$ and $\boldsymbol{E}$ and $\boldsymbol{X}$ are independent. Provided $\boldsymbol{I}_p - \boldsymbol{A}$ is invertible where $\boldsymbol{I}_p$ is a $p \times p$ identity matrix, model (1) can be equivalently expressed as

$$\begin{aligned} \boldsymbol{Y} \mid \boldsymbol{X} &\sim \mathrm{N}_p\left\{(\boldsymbol{I}_p - \boldsymbol{A})^{-1}\boldsymbol{B}\boldsymbol{X}, (\boldsymbol{I}_p - \boldsymbol{A})^{-1}\boldsymbol{\Sigma}(\boldsymbol{I}_p - \boldsymbol{A})^{-T}\right\}, & (2) \\ \boldsymbol{X} &\sim \mathrm{N}_{2p}(\boldsymbol{0}, \boldsymbol{\Psi}). & (3) \end{aligned}$$

To link an SEM to an RGM, we draw a path diagram $\mathcal{G} = (V, E)$ of SEM by the following rules:

(i) define vertices $V = \{1, \ldots, p, p+1, \ldots, 3p\}$ which represent $(\boldsymbol{Y}, \boldsymbol{X}) = (Y_1, \ldots, Y_p, X_1, \ldots, X_{2p})$;

(ii) draw directed edges $E^d = \{(j, i) \mid a_{ij} \neq 0 \text{ or } b_{i,j-p} \neq 0\}$; and

(iii) draw undirected edges $E^u\{\{i, j\} \mid \boldsymbol{\Psi}_{i-p,j-p} \neq 0\}$.

In words, (i) we introduce a node for each variable in $(\boldsymbol{Y}, \boldsymbol{X})$ with nodes $j = 1, \ldots, p$ corresponding to Y_j and $p + j$ corresponding to X_j for $j = 1, \ldots, 2p$; (ii) nodes $i = 1, \ldots, p$ (i.e. Y_i nodes) become targets of directed edges from node j if the corresponding $a_{ij} \neq 0$ or $b_{i,j-p} \neq 0$; (iii) we introduce undirected edges between X_i and X_j (i.e. nodes $i + p$ and $j + p$) if $\boldsymbol{\Psi}_{i,j} \neq 0$. Let $A = \mathrm{diag}(A_1, A_2)$ denote a block diagonal matrix with diagonal blocks A_1 and A_2. Figure 2 shows an example of an RGM with $p = 2$,

$$\boldsymbol{A} = \begin{bmatrix} 0 & * \\ * & 0 \end{bmatrix}, \ \boldsymbol{B} = \begin{bmatrix} * & * & 0 & 0 \\ 0 & 0 & * & 0 \end{bmatrix} \text{ and } \boldsymbol{\Psi} = \mathrm{diag}\left(\begin{bmatrix} * & * \\ * & * \end{bmatrix}, \boldsymbol{I}_2\right),$$

with $*$ indicating non-zero elements. In general, the path diagram of an SEM following rules (i)-(iii) is not always an RGM. However, the probability distribution of $(\boldsymbol{Y}, \boldsymbol{X})$ defined

by a SEM is Markov with respect to the path diagram $\mathcal{G}$ (Spirtes, 1995; Koster, 1996) if $\boldsymbol{\Sigma}$ is diagonal $\boldsymbol{\Sigma} = \text{diag}(\sigma_1, \ldots, \sigma_p)$ and $\boldsymbol{\Psi}$ is block diagonal with each block being a full matrix (without zeros). Since the main inferential goal of this paper is to investigate the regulatory relationship between genes, we will not model the marginal distribution of $\boldsymbol{X}$ and will only focus on the conditional distribution of $\boldsymbol{Y} \mid \boldsymbol{X}$. We fix b_{ij} to 0 for $j \neq 2i-1$ or $2i$ because copy number and methylation of gene i, in principle, only affect the expression of gene i.

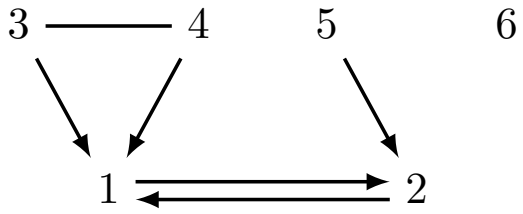

Figure 2: Path diagram for simultaneous equations.

2.4 Priors and graph structure learning

The structural zeros in $\boldsymbol{A}$ and $\boldsymbol{B}$ correspond to missing edges in the RG. Therefore, learning the graph structure is equivalent to finding sparse estimators for $\boldsymbol{A}$ and $\boldsymbol{B}$. Towards this end, we define a non-local prior for $\boldsymbol{A}$ and $\boldsymbol{B}$, which is constructed as follows. We put a thresholded prior on each element of $\boldsymbol{A}$ and $\boldsymbol{B}$. We first write a_{ij} and b_{ik} as

$$a_{ij} = \tilde{a}_{ij} \text{I}(\mid \tilde{a}_{ij} \mid > t_i) \quad \text{and} \quad b_{ik} = \tilde{b}_{ik} \text{I}(\mid \tilde{b}_{ik} \mid > t_i),$$

for $i = 1, \ldots, p$, $j \neq i$ and $k = 2i-1, 2i$. The threshold parameter t_i controls the minimum effect sizes of a_{ij} and b_{ik}. We do not fix t_i but instead learn it from the data by assigning a prior $t_i \sim \text{Uniform}(0, t_0)$. We impose normal priors for $\tilde{a}_{ij}$ and $\tilde{b}_{ik}$, $\tilde{a}_{ij} \sim \text{N}(0, \tau_{ij})$ and $\tilde{b}_{ik} \sim \text{N}(0, \nu_{ik})$ and conjugate hyperpriors $\tau_{ij} \sim \text{IG}(\alpha_\tau, \beta_\tau)$ and $\nu_{ik} \sim \text{IG}(\alpha_\nu, \beta_\nu)$. Ni et al. (2016) show that marginally, after integrating with respect to t_i, the induced marginal prior for a_{ij} or b_{ik} is a mixture of a point mass at 0 and a non-local prior (Johnson and Rossell, 2010). Non-local priors shrink small effects to zero, which is desirable in our setting where

we are only interested in edges with moderate to strong effects. We complete our prior specifications with a conjugate prior $\sigma_i \sim \text{IG}(\alpha_\sigma, \beta_\sigma)$.

Posterior inference is straightforward by Markov chain Monte Carlo (MCMC). It is worth mentioning that we do not need to check the invertibility of $\boldsymbol{I} - \boldsymbol{A}$ because $\boldsymbol{I} - \boldsymbol{A}$ is practically always invertible and we never have to invert the matrix in our MCMC. Details are provided in Supplementary Material A. Let $\boldsymbol{\Gamma}_A = (\gamma_{A,ij})$ and $\boldsymbol{\Gamma}_B = (\gamma_{B,ik})$ denote matrices indicating non-zero elements in $\boldsymbol{A}$ and $\boldsymbol{B}$, respectively. An estimated graph can be reported by selecting edges for which the marginal posterior probability $p(\gamma_{A,ij} = 1 \mid \text{Data})$ or $p(\gamma_{B,ik} = 1 \mid \text{Data})$ exceeds a certain cutoff. This is the median probability model (MPM) when the cutoff is fixed at 1/2. Instead of being fixed at an arbitrary value, the cutoff can be chosen to control the posterior expected false discovery rate (FDR) at a desired level α (Newton et al., 2004; Müller et al., 2006). Alternatively we report the highest posterior probability model (HPM) by maximizing $p(\boldsymbol{\Gamma}_A, \boldsymbol{\Gamma}_B \mid \text{Data})$. However, this is only feasible when the model space is small (i.e. with very small p).

3 Simulations

We carry out a simulation study to validate the model's ability to recover a true graph with sample sizes similar to the later application. We generate data by mimicking the actual data from the later application, using a sample size of $n = 276$ and $p = 10$ genes. For each gene, we generate two hypothetical DNA level measurements $\boldsymbol{X}$ from (3) with $\boldsymbol{\Psi} = \boldsymbol{I}_{2p}$ (corresponding to copy number and methylation in the real data).

We consider three scenarios. In scenario 1, we randomly set 4/5 entires in $\boldsymbol{A}$ and 1/3 entries in $\boldsymbol{B}$ to zero. To ensure that the true model is identifiable, we restrict the simulation truth to include at least one edge from DNA level measurements to each gene (this is relaxed in scenario 2). We draw nonzero elements of $\boldsymbol{A}$ and $\boldsymbol{B}$ from $\{-0.5, 0.5\}$ with equal probability and let $\boldsymbol{\Sigma} = 0.25\boldsymbol{I}_{10}$ in the simulation truth.

In scenario 2, we reduce the signal by drawing nonzero elements of $\boldsymbol{A}$ and $\boldsymbol{B}$ from $\{-0.4, 0.4\}$ and increase the noise to $\boldsymbol{\Sigma} = \boldsymbol{I}_{10}$. Moreover, we do not force every gene to be connected to at least one of its DNA level measurements. As a result one gene is independent of both its DNA level measurements. For scenarios 1 and 2, $\boldsymbol{Y}$ is then

generated from (2).

Scenario 3 is the same as scenario 1, except that the conditional distribution of $\boldsymbol{Y} \mid \boldsymbol{X}$ is misspecified and generated from a p-dimensional multivariate t-distribution $\mathrm{T}_p(\boldsymbol{\mu}, \boldsymbol{\Theta}, \delta)$ with location $\boldsymbol{\mu} = (\boldsymbol{I}_p - \boldsymbol{A})^{-1}\boldsymbol{B}\boldsymbol{X}$, scale matrix $\boldsymbol{\Theta} = (\boldsymbol{I}_p - \boldsymbol{A})^{-1}\boldsymbol{\Sigma}(\boldsymbol{I}_p - \boldsymbol{A})^{-T}$ and degrees of freedom $\delta = 3$. By design, scenarios 2 and 3 are more difficult and perhaps more realistic than scenario 1.

The hyperparameters are specified as $\alpha_\sigma = \beta_\sigma = \alpha_\tau = \beta_\tau = \alpha_\nu = \beta_\nu = 0.01$ and $t_0 = 1$. We run 50,000 MCMC iterations, discard the first 25,000 iterations as burn-in and retain every 5th sample. The graph is estimated as an MPM.

In Table 1 we report the true positive rate (TPR), FDR and Matthews correlation coefficient (MCC). We also plot the receiving operating characteristic (ROC) curves in Figure 3. The ROC curve, TPR and FDR are for classifying entries in $\boldsymbol{A}$ and $\boldsymbol{B}$ as non-zero ("positive") or not ("negative"). That is, for the inclusion of all possible edges. TPR and FDR are the realized error rates under the reported MPM graph. The area under the ROC curve (AUC) is reported in Table 1 (the ROC plot for scenario 1 is not shown – the AUC is 1). The performance in scenario 1 is nearly perfect. As expected, scenarios 2 and 3 are more challenging. In particular, we find a higher FDR, due to the lower signal-to-noise ratio and a minor identifiability issue for scenario 2 and model misspecification for scenario 3.

Table 1: Simulation results. The average operating characteristics are calculated on the basis of 50 simulations; standard deviations are given within parentheses.

Scenario	MCC	TPR	FDR	AUC
1	0.99 (0.02)	1.00 (0.00)	0.01 (0.02)	1.00 (0.00)
2	0.76 (0.13)	0.98 (0.03)	0.26 (0.14)	0.97 (0.04)
3	0.87 (0.08)	0.97 (0.04)	0.14 (0.08)	0.99 (0.02)

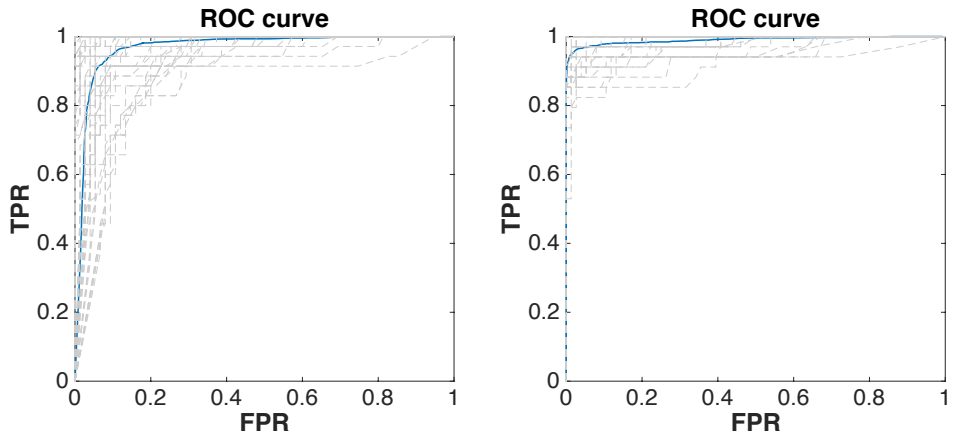

(a) Scenario 2 (b) Scenario 3

Figure 3: ROC curves for scenarios 2 (panel a) and 3 (panel b). Each dashed line is an ROC curve for one simulation and the solid line is the average ROC curve over 50 simulations.

4 Application

4.1 ZODIAC pairwise gene interaction analysis

We use TCGA-Assembler (Zhu et al., 2014) to retrieve mRNA gene expression (GE), DNA copy number (CN) and DNA methylation (ME) data for different cancer types from TCGA. For illustration, we report the analysis based on pairs of genes whose GEs are moderately correlated (between 0.5 and 0.7) and show high correlations (> 0.8) with corresponding CNs or MEs. The resulting dataset consists of 1395 pairs from 15 cancers (details of the data are in secton B of the Supplementary Material). We apply our method separately to each pair of genes using the same hyperparameter and MCMC settings as in the simulation study. Since $p = 2$ is sufficiently small, we can report the HPM as an estimate of the graph. In this application, we aim to demonstrate: (1) like ZODIAC, our model can be used as a hypothesis generating tool; (2) compared to the undirected graph model used in ZODIAC, the RGM is more flexible and allows for more and different inference and interpretation.

Among the 1395 gene pairs, 518 pairs are estimated to have one-way interactions, 131 pairs have two-way interactions and the rest have no interactions. In addition, we find 2544 intragenic GE-CN interactions and 1157 intragenic GE-ME interactions. We show three gene pairs in Figure 4. These interactions are confirmed by ZODIAC (Zhu et al., 2015) but our approach has richer interpretations. In Figure 4a, both our method and ZODIAC

find a negative interaction between MEA1 and UBR2 whereas our analysis indicates that MEA1 downregulates UBR2 specifically in rectum adenocarcinoma (READ). Similarly, in Figure 4b, while ZODIAC finds positive interaction between PHTF2 and TLK1, our model suggests a positive feedback loop between these two genes in testicular germ cell tumors (TGCT). Figure 4c manifests a unique feature of our model, a negative feedback loop between COG3 and MRPS31 in colon adenocarcinoma (COAD). Negative feedback loops are a common and important mechanism in gene regulations. For example, it can reduce expression noise and intercellular variability in molecular levels (Singh, 2011). But negative feedback loops are not allowed by DAGs and MRFs, both of which include the simplified assumption that gene interaction is always symmetric, either positive or negative. Our model is able to generate hypotheses that involve both one-way and two-way interactions and both positive and negative feedback loops.

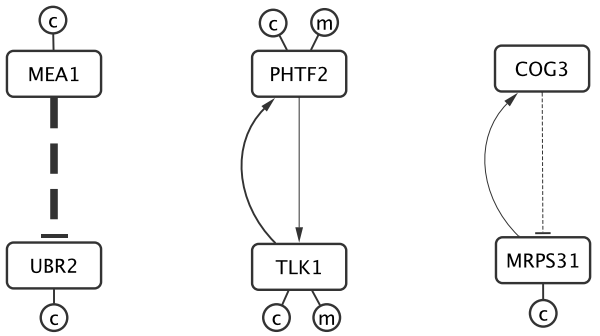

(a) READ (b) TGCT (c) COAD

Figure 4: Interactions of three pairs of genes in three cancers. Genes attached by tags "c" and "m" indicate that they are associated with copy number and methylation, respectively. Solid lines with arrowheads represent stimulatory interactions, whereas dashed lines with horizontal bars denote inhibitory regulation. Edge width is proportional to its posterior probability.

4.2 Colon adenocarcinoma pathway analysis

We analyze data from colon adenocarcinoma (COAD) samples. We focus on genes that are mapped to the RAS-MAPK pathway, which is critical for initiation of carcinogenesis in COAD (Colussi et al., 2013). The RAS-MAPK pathway includes $p = 10$ core genes. Restricting to samples with available GE, CN and ME data, the sample size is $n = 276$. We

run two parallel MCMC simulations, each with 50,000 iterations, discard the first 50% as burn-in and thin the chains to every 5th sample. MCMC diagnostics show no evidence for lack of practical convergence (for details see part C of the Supplementary Materials). We summarize the posterior distribution on the unknown graph by controlling the posterior expected FDR $< 10\%$.

We find that all genes are associated with their respective copy number and NRAS, MAP2K1, MAPK1 and MAPK3 are associated with methylation as well. Therefore, the gene interactions are fully identifiable by the argument in the last paragraph of Section 2.2. The full network of gene interactions is shown in Figure 5a. The solid lines with arrowheads represent stimulatory interactions and the dashed lines with perpendicular bars denote inhibitory regulations. Uncertainties about the selected edges are quantified as posterior inclusion probabilities which are displayed as edge width in the figure. In total, we find 21 stimulatory and 7 inhibitory regulatory relationships.

Gene networks are often made up of a small set of recurrent regulation patterns, called network motifs, which can be thought of as fundamental building blocks for the network and are expected to occur more often in gene networks than in random networks. In Figures (5b)-(5e), we display four motifs identified by our method that are commonly observed in gene networks (Alon, 2007). Figure (5b) shows a feed-forward loop among SOS2, KRAS and MAPK3. Part of this feed-forward loop is well studied in COAD (Zenonos and Kyprianou, 2013). SOS2 binds KRAS and removes guanosine diphosphate (GDP) molecules from KRAS and thus allows guanosine triphosphate (GTP) molecules to bind and activate it. The active KRAS would eventually activates MAPK3 through the kinase cascade. Another important motif is feedback loop. We present a negative and a positive feedback loop in Figures 5c and 5d, respectively. The regulatory relationships in the negative feedback loop of SOS2, KRAS and MAP2K1 have been studied extensively. SOS2$\rightarrow$KRAS$\rightarrow$MAP2K1 is part of the well-known MAP kinases cascade (Plotnikov et al., 2011; Zenonos and Kyprianou, 2013) while MAP2K1 can phosphorylate and inhibit SOS2 and thereby reduces MAP2K1 activation (Holt et al., 1996; Mendoza et al., 2011). For the positive feedback loop (Figure 5d), MAP2K1 activating MAPK1 is again part of the MAP kinases cascade but the reversed activation is less explored in the literature. Similarly,

another network motif (Figure 5e), regulatory cascade, also need to be validated by further biological experiment.

On the gene level, we calculate the posterior distribution of degree of each gene, that is, the number of edges connected to the gene. We visualize these posterior distributions by box plots in Figure 6a. Highly connected genes are often called hub genes which are usually more involved in multiple regulatory activities than non-hub genes. MAPK1 and MAPK3 appear to be the two most highly connected genes. Several studies have shown that overexpression of MAPK1/3 plays an critical role in the progression of COAD (Fang and Richardson, 2005) and is responsible for proliferation, differentiation, survival, migration and angiogenesis of tumors in many cancers (Dhillon et al., 2007).

For another quantitative exploration of the known MAPK signalling cascade we consider the following inference. The following signaling pathway/cascade has been extensively studied and validated in biological literatue (Plotnikov et al., 2011; Zenonos and Kyprianou, 2013),

$$\left\{\begin{array}{c}\text{GRB2}\\ \text{SOS1}\\ \text{SOS2}\end{array}\right\} \longrightarrow \left\{\begin{array}{c}\text{NRAS}\\ \text{KRAS}\end{array}\right\} \longrightarrow \text{BRAF} \longrightarrow \left\{\begin{array}{c}\text{MAP2K1}\\ \text{MAP2K2}\end{array}\right\} \longrightarrow \left\{\begin{array}{c}\text{MAPK1}\\ \text{MAPK3}\end{array}\right\} \tag{4}$$

where genes within each curly brace belong to the same gene family (except for GRB2 for which the protein often binds to SOS to form a protein complex) and play similar roles in the pathway. We label each gene in cascade (4) from left to right with integer $j = 1, \ldots, p$ as our reference ordering. The labels within each curly brace are arbitrary. To compare our findings with this well established cascade, we score each gene j by

$$\text{Score}_j = \text{indegree}_j - \text{outdegree}_j, \tag{5}$$

where indegree_j is the number of edges pointing towards gene j and outdegree_j is the number of edges pointing away from gene j. Intuitively, gene j is likely to be on the left of gene k in cascade (4) if $\text{Score}_j < \text{Score}_k$. We then rank the scores in an increasing

order and plot the rank against the reference ordering in Figure 6b. We also calculate the normalized Kendall's tau distance, which is the ratio of the number of discordant pairs over the total number of pairs. The normalized Kendall's tau distance lies between 0 and 1, with 0 indicating a perfect agreement of the two orderings and 1 indicating a perfect disagreement. The resulting Kendall's tau distance is 0.07. Our findings appear to be consistent with the biologically validated pathway since both Figure 6b and the Kendall's tau distance indicate a good concordance between our rankings and the reference ordering.

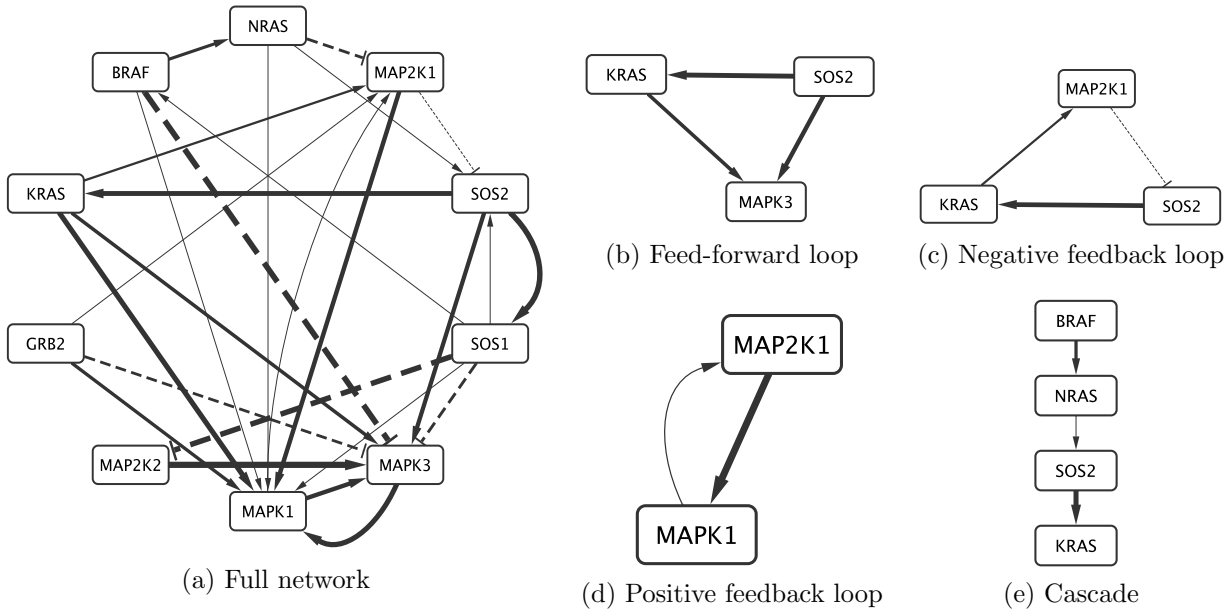

(a) Full network (b) Feed-forward loop (c) Negative feedback loop (d) Positive feedback loop (e) Cascade

Figure 5: Estimated gene network for colon adenocarcinoma. Solid lines with arrowheads represent stimulatory interactions, whereas dashed lines with horizontal bars denote inhibitory regulation. Line width is proportional to its posterior probability. Panel (a) full network; Panels (b)-(e) network motifs.

5 Discussion

In this article, we have introduced a Gaussian RGM to model gene regulatory relationships from genomic data. RGMs are statistically more general and biologically more interpretable than MRFs and DAGs. By integrating DNA level information, we are able to differentiate between RGMs that belong to the same Markov equivalence class. We exploit the connection between RGMs and SEMs for efficient inference. We constructed a prior probability

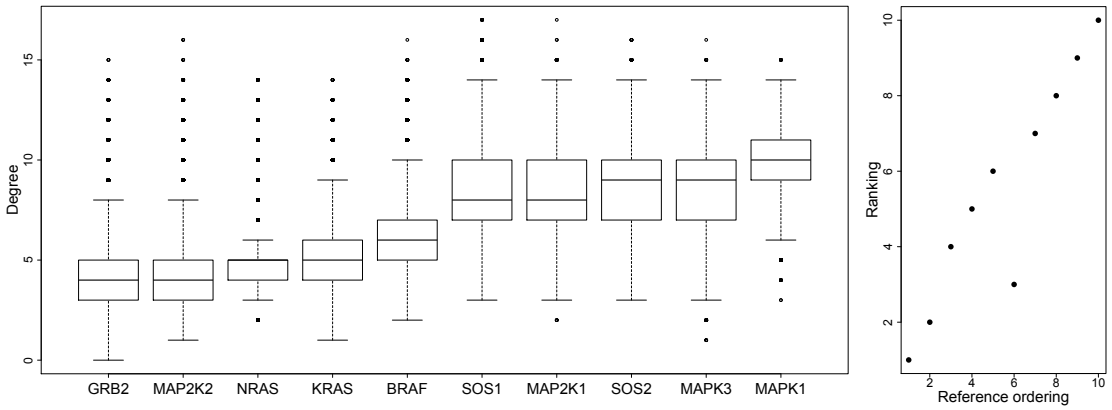

(a) (b)

Figure 6: Panel (a): box plots for posterior distribution of degree of each gene; Panel (b): scatter plot for ranking vs reference ordering. Ranking of the genes is based on the Score_j defined by (5). The reference ordering is based on (4) where the genes with the same bracket are assigned arbitrary ordering.

model for the unknown graph using a thresholded model which marginally defines a mixture of non-local prior and a point mass. We use simulation studies to illustrate the performance of our method in terms of graph structure learning. Our method is applied to a ZODIAC gene interaction analysis and a colon cancer pathway analysis. Some of our findings are consistent with the literature, while others need to be validated by biological experiments. Although our applications focus on GRNs, the proposed approach is general and can be potentially applied to other scientific settings such as climate sciences and macroeconomics. The approach works for any network where some edges have known direction if included (inclusion itself need not be fixed).

The link between RGMs and SEMs is based on the assumption that the gene expressions are multivariate Gaussian, which could be thought of as one limitation of our model. We have empirically demonstrated the robustness of the proposed model against slight model misspecification. More general sampling models might need an additional hierarchical layer of latent variables, such as latent probit scores for binary outcomes.

References

Alon, U. (2007). Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6):450–461.

Colussi, D., Brandi, G., Bazzoli, F., and Ricciardiello, L. (2013). Molecular pathways involved in colorectal cancer: implications for disease behavior and prevention. *International journal of molecular sciences*, 14(8):16365–16385.

Dhillon, A. S., Hagan, S., Rath, O., and Kolch, W. (2007). Map kinase signalling pathways in cancer. *Oncogene*, 26(22):3279–3290.

Dobra, A., Lenkoski, A., and Rodriguez, A. (2012). Bayesian inference for general gaussian graphical models with application to multivariate lattice data. *Journal of the American Statistical Association*.

Fang, J. Y. and Richardson, B. C. (2005). The mapk signalling pathways and colorectal cancer. *The lancet oncology*, 6(5):322–327.

Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airoldi, E. M. (2010). A survey of statistical network models. *Foundations and Trends® in Machine Learning*, 2(2):129–233.

Green, P. J. and Thomas, A. (2013). Sampling decomposable graphs using a markov chain on junction trees. *Biometrika*, 100(1):91–110.

Hoff, P. (2009). Multiplicative latent factor models for description and prediction of social networks. *Computational & Mathematical Organization Theory*, 15(4):261–272.

Holt, K. H., Kasson, B. G., and Pessin, J. E. (1996). Insulin stimulation of a mek-dependent but erk-independent sos protein kinase. *Molecular and Cellular Biology*, 16(2):577–583.

Johnson, V. E. and Rossell, D. (2010). On the use of non-local prior densities in bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(2):143–170.

Koster, J. T. (1996). Markov properties of nonrecursive causal models. *The Annals of Statistics*, 24(5):2148–2177.

Mendoza, M. C., Er, E. E., and Blenis, J. (2011). The ras-erk and pi3k-mtor pathways: cross-talk and compensation. *Trends in biochemical sciences*, 36(6):320–328.

Mitra, R., Müller, P., Liang, S., Yue, L., and Ji, Y. (2013). A bayesian graphical model for chip-seq data on histone modifications. *Journal of the American Statistical Association*, 108(501):69–80.

Müller, P., Parmigiani, G., and Rice, K. (2006). Fdr and bayesian multiple comparisons rules.

Newton, M. A., Noueiry, A., Sarkar, D., and Ahlquist, P. (2004). Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, 5(2):155–176.

Ni, Y., Stingo, F., and Baladandayuthapani, V. (2016). Bayesian graphical regression. *submitted*.

Ni, Y., Stingo, F. C., and Baladandayuthapani, V. (2015). Bayesian nonlinear model selection for gene regulatory networks. *Biometrics*, 71(3):585–595.

Plotnikov, A., Zehorai, E., Procaccia, S., and Seger, R. (2011). The mapk cascades: signaling components, nuclear roles and mechanisms of nuclear translocation. *Biochimica et Biophysica Acta (BBA)-Molecular Cell Research*, 1813(9):1619–1633.

Shin, S.-Y., Rath, O., Zebisch, A., Choo, S.-M., Kolch, W., and Cho, K.-H. (2010). Functional roles of multiple feedback loops in erk and wnt signaling pathways that regulate epithelial-mesenchymal transition. *Cancer research*, 70(17):6715.

Singh, A. (2011). Negative feedback through mrna provides the best control of gene-expression noise. *NanoBioscience, IEEE Transactions on*, 10(3):194–200.

Spirtes, P. (1995). Directed cyclic graphical representations of feedback models. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 491–498. Morgan Kaufmann Publishers Inc.

Stingo, F. C., Chen, Y. A., Vannucci, M., Barrier, M., and Mirkes, P. E. (2010). A bayesian graphical modeling approach to microrna regulatory network inference. *The annals of applied statistics*, 4(4):2024.

TCGA (2012). Comprehensive molecular characterization of human colon and rectal cancer. *Nature*, 487(7407):330–337.

Telesca, D., Müller, P., Kornblau, S. M., Suchard, M. A., and Ji, Y. (2012a). Modeling protein expression and protein signaling pathways. *Journal of the American Statistical Association*, 107(500):1372–1384.

Telesca, D., Müller, P., Parmigiani, G., and Freedman, R. S. (2012b). Modeling dependent gene expression. *The Annals of Applied Statistics*, 6(2):542–560.

Wang, H. and West, M. (2009). Bayesian analysis of matrix normal graphical models. *Biometrika*, 96(4):821–834.

Wang, W., Baladandayuthapani, V., Holmes, C. C., and Do, K.-A. (2013). Integrative network-based bayesian analysis of diverse genomics data. *BMC Bioinformatics*, 14(Suppl 13):S8.

Yajima, M., Telesca, D., Ji, Y., and Müller, P. (2015). Detecting differential patterns of interaction in molecular pathways. *Biostatistics*, 16(2):240–251.

Zenonos, K. and Kyprianou, K. (2013). Ras signaling pathways, mutations and their role in colorectal cancer. *World J Gastrointest Oncol*, 5(5):97–101.

Zhang, D., Wells, M. T., Turnbull, B. W., Sparrow, D., and Cassano, P. A. (2005). Hierarchical graphical models: An application to pulmonary function and cholesterol levels in the normative aging study. *Journal of the American Statistical Association*, 100(471):719–727.

Zhu, Y., Qiu, P., and Ji, Y. (2014). Tcga-assembler: open-source software for retrieving and processing tcga data. *Nature methods*, 11(6):599–600.

Zhu, Y., Xu, Y., Helseth, D. L., Gulukota, K., Yang, S., Pesce, L. L., Mitra, R., Müller, P., Sengupta, S., Guo, W., et al. (2015). Zodiac: a comprehensive depiction of genetic interactions in cancer by integrating tcga data. *Journal of the National Cancer Institute*, 107(8):djv129.